

0/890869

PCT/NZ00/00012

Intellectual
Property Office

of New Zealand

Te Pou Rāhui Hanga Hou



4

REC'D 07 APR 2000

WIPO PCT

CERTIFICATE

This certificate is issued in support of an application for Patent registration in a country outside New Zealand pursuant to the Patents Act 1953 and the Regulations thereunder.

I hereby certify that annexed is a true copy of the Complete Specification as filed on 11 February 1999 with an application for Letters Patent number 334139 made by PEC (NEW ZEALAND) LIMITED.


I further certify that pursuant to a claim filed on 9 February 1999 under Section 24(1) of the Patents Act 1953, a direction that the application proceed in the name of CARDAX INTERNATIONAL LIMITED by virtue of a deed dated 1 June 1999.

Dated 10 March 2000.

**PRIORITY
DOCUMENT**
SUBMITTED OR TRANSMITTED IN
COMPLIANCE WITH RULE 17.1(a) OR (b)



Neville Harris
Commissioner of Patents



NEW ZEALAND

Patents Act 1953


PROVISIONAL SPECIFICATION

IMPROVEMENTS IN OR RELATING TO
CONTROL AND/OR MONITORING SYSTEMS

We, PEC (NEW ZEALAND) LIMITED a New Zealand company of 2
Station Road, Marton, New Zealand do hereby declare this
invention to be described in the following statement:-

334139

-1-



This invention relates to improvements in or relating to control and/or monitoring systems and a surveillance camera for use with such systems.

In a typical central access control and/or monitoring system of, say, a large commercial or industrial complex or a plurality of different sites, the person(s) at the front end of the system has access to a considerable volume of information. This can typically include a multiplicity of screens which must continually be monitored. Effectively, what can occur is an "information overload" where the person(s) at the front end of the system cannot continually assimilate all of the available information in a manner which facilitates an efficient/effective monitoring function. As a result "alarm" or "unusual" situations appearing on a screen or number of screens can be overlooked or not acted upon as quickly as desired.

There is, therefore, a need for a monitoring system whereby information relating to an unusual situation or an alarm situation can be distinguished from routine information thereby increasing the likelihood of the information concerning the unusual/alarm situation being noted and acted upon more expeditiously.

Accordingly, it is one object of the present invention to provide a surveillance camera which is able to perform data processing.

A further object of the invention is to provide a monitoring system wherein data processing can occur at an image capture unit.

Broadly, according to one aspect of the invention, there is provided a surveillance camera comprising a digital video camera with integrated intelligence.

According to a further broad aspect of the invention there is provided a monitoring system including at least one surveillance camera according to the first broad aspect.

According to yet a further broad aspect of the invention there is provided a method of monitoring including the steps of capturing image data at an image capture unit, effecting at least some data processing at the image capture unit and causing selected data to be transferred from the image capture unit.

In the preferred method the image capture unit is a digital video camera with integrated intelligence.

In the following more detailed description of preferred embodiments of the invention reference will be made to the accompanying drawings in which:-

Figure 1 is a schematic diagram of the functional blocks of the architecture of the digital video camera,

Figure 2 is a schematic diagram of the camera imaging hardware,

Figure 3 is a schematic illustration of the image buffer memory organisation,

Figure 4 is a state transition diagram for an image,

Figures 5 to 7 are circuit diagrams relating to one embodiment of the digital video camera, and

Figure 8 is a schematic of an access control system incorporating the digital video camera.

To more fully describe the invention according to its various aspects reference will be made to a digital video camera (herein "DVC") according to the invention. This camera can be employed in a monitoring system (including access control) an example of which is our CARDAX system. The present disclosure will thus, by way of example, refer to a preferred embodiment of the DVC as forming

part of an enhanced version of our CARDAX access control and monitoring system.

According to the present invention, the DVC has integrated intelligence such that the camera itself can effect data processing. Data processing is thus performed on captured images at the camera thereby enabling the camera to rapidly analyse image data and perform functions which can include motion detection, motion detection within a region of the camera's field of view (the region being, eg specified by a system operator), false detection recognition (ie curtains moving in a breeze) and tracking objects to record the path of movement through the field of view (useful for eg detecting access to an area in the camera's field of view in authorised and unauthorised directions.

A further ability of the DVC is to indelibly stamp images with a digital "watermark" to authenticate images. It is proposed that such watermarking would be able to be used as Court evidence so that in connection with an image linked to an event the watermark can verify which camera (in a system) captured the image, the date and time the image was captured and verifying that the image has not been tampered with in any way since it was captured.

Used in, for example, an enhanced version of CARDAX the system architecture will include one or more DVC's connected to the "Local BUS". Compressed digital data

can then be sent on request from the DVC over the Local BUS to the CARDAX controller and then over a high speed communications medium to the host system. This host system can be based on a personal computer as a central processing point for a system set-up, monitoring and reporting. The CARDAX door controllers will monitor and control access to doors. Each of the door controllers can interface to a variety of devices via the Local BUS RS485 link. Devices on the Local BUS can include door readers, door strike controllers and input/output units.

In its simplest application the DVC can be used both as a surveillance camera and as a replacement for a standard passive infra-red motion detector. When performing motion detection the camera will be able to generate an event to be logged by the system host in the same manner as a contact alarm for a traditional passive infra-red detector. Under software control from an operator workstation the DVC can provide images on screen with the operator being able to select any camera in the system.

By virtue of its "intelligence" the DVC can analyse images captured and be able to process the images for subsequent transmission to the host system or storage for later transmission of any images arising from motion detection, false detection or tracking objects. To indicate an unusual or alarm situation the camera can immediately on-pass such information to the host system for analysis at the operator work-station. In this

manner, the camera can distinguish between routine information and alarm/unusual situation information thereby resulting in the operator work-station only being provided with information on which immediate analysis/action is required. An alarm function can be included to alert the operator to the alarm/unusual situation so as to ensure that the images transmitted from the digital camera are immediately analysed.

Referring firstly to Figure 1 of the drawings the DVC consists of six basic functional blocks as shown. The communications block 10 is in the system being described herein the CARDAX Local BUS interface for the camera. As will be appreciated by those skilled in the art galvanic isolation can be provided between the camera electronics and the Local BUS communication lines which can also be provided with protective devices.

The communications block 10 is based around a microprocessor with various integrated peripherals including parallel and serial ports. The microprocessor is provided with FLASH E² PROM non-volatile memory to hold programme code and set-up data and volatile SRAM memory for general use.

The non volatile memory holds programme code and set-up data for both the communications block 10 and the main processing block 11. The communications block 10 downloads main programme code and set-up data as required

(for example, at power-up reset) to the main processing block 11 which is provided only with volatile memory for data and programme control. In normal operation programme code and set-up data can be updated/down-loaded from an external device via the Local BUS interface. Several different versions of programme code and set-up data for both the communications block 10 and the main processing block 11 may be stored in the non-volatile memory. Software control mechanisms select the version(s) to be used in any particular time.

The communications block 10 additionally includes a general purpose serial communications interface (RS232 port) which can be used to connect, for example, a lap-top computer. This port would typically be used for field diagnostics and is preferably only accessible physically when a cover is removed from the camera.

The communications block 10 can also provide unit address initialisation, unique electronic serial number identification, tamper detection, micro-controller supervisory functions, and visual processing indications using LEDs.

The main processing block 11 is based around a digital signal processor ("dsp"). The dsp is provided with SRAM which it can flexibly segment into areas of data and programme code memory. The dsp primarily analyses images and compresses data.

A multi-functional "host port" associated with the dsp in processing block 11 connects to the address and data BUS system of the communications block 10 to provide a high-speed bi-directional, parallel data transfer link between the two functional blocks. The dsp appears to the microprocessor in the communications block 10 as a local peripheral in its I/O map.

An image capture block 12 is based around a CMOS active pixel sensor (APS) imaging chip which is effectively a camera-on-a-chip and the necessary bias circuitry for it. The image capture block 12 acquires images and presents them as a stream of digital data, on an 8-bit wide data BUS to the video buffer block 13. An image from the video buffer block can vary in size, eg it can be a full frame or a window (rectangular block) from within the full frame.

Infra-red LEDs can be included in the image capture block 12 to provide infra-red illumination. The DVC can, therefore, monitor either in the visual light spectrum or in the infra-red spectrum.

The video buffer block 13 consists of SRAM for storing image data acquired by the image capture block 12. While another image is being accumulated in the image capture block 12, the main processing block 11 has access to the data stored in the video buffer block 13.

A set of BUS transceivers interface the SRAM in the video buffer block 13 to the 8-bit data BUS system of the image capture block 12 and another set of BUS transceivers interface the SRAM to the 24-bit data BUS system of the main processing block 11. Through the transceivers, dual port access to the SRAM is provided but only one "port" can be accessed at a time. The APS in the image capture block can effectively write 8-bit data into the SRAM of the video buffer block 13.

The dsp can effectively read data from and write data to the SRAM of the video buffer block 13 as 24-bit words.

Image data from the image capture block 12 is passed to the video buffer block 13, 1-BYTE (8-bit) at a time. As each BYTE is passed over it is loaded into one of three BUS transceivers such that every three consecutive BYTES are loaded to effectively form a 3-BYTE (24-bit) word. When 3 BYTES have accumulated in the BUS transceivers they are simultaneously latched to the transceiver outputs and then written into the SRAM as a single 24-bit word.

A sensor/buffer control block 14 consists of a small microprocessor and various discrete logic. Much of the discrete logic is in a programmable logic device (pld). A full-duplex, asynchronous serial link between the dsp and the image capture microprocessor results in the dsp

being able to request a variety of options that the microprocessor in the image capture or sensor/buffer control block 14 and the APS are able to provide. The infra-red LEDs are controlled from the microprocessor in the control block 14. The microprocessor in block 14 communicates with the sensor over an I²C BUS synchronous serial link.

The power supply block 15 takes power from an external source and generates all the necessary power rails for the DVC electronics. Line protection and filter circuitry are included.

In a preferred form of the invention as herein described the following power rails are generated:-

5V for digital circuits in the communications block 10 and image capture block 12;

5V for analogue circuits in the image capture block 12;

5V programming supply for the non-volatile memory in the communications block 10;

3.3V for digital circuits in the main processing, video buffer and sensor-buffer control blocks 11 and 14;

7-10V for the IR LEDs.

Typically, an image captured by the APS is transferred directly into the video buffer 13 via one BUS interface. The dsp then accesses this by another BUS interface and processes it. Compressed images can be stored until they are transferred from the main processing block 11 to the communications block 10 which can transmit them over the Local BUS or overwritten by newer images.

The electronics for the DVC resides on two multi-layer printed circuit board (pcb) assemblies. One is a small assembly consisting of the APS and the bias circuitry and oscillator associated with it, the lens for the APS and IR LEDs used for illumination in dark environments. All other electronics reside on the second larger pcb. The two assemblies are electrically inter-connected using a flexible pcb strip. The pcbs and lens assemblies are housed within a sealed plastic enclosure. A plastic mounting bracket preferably secures the small image capture pcb assembly to the larger pcb assembly in a manner that allows the APS, lens and IR LEDs to be rotated in one plane. This allows a degree of freedom of movement for the module to be aligned during installation.

According to a preferred form the mounting is such that it is possible to mount the DVC at any angle on (for example, but not restricted to) a wall, a ceiling, or an

adjustable bracket that allows the camera to be easily pointed at an appropriate view. The standard combination of a wide-angle lens and having the image capture assembly internally mounted at a 45° angle means that in most cases the DVC can simply be placed at a location without any internal adjustment being required. In cases where internal adjustment of the image capture assembly is required, it can be rotated up to 45°.

Installing the DVC according to a preferred form involves a plastic mounting plate. The plate is affixed to a wall, ceiling, pan/tilt mount or other camera mount and provides terminals for attaching external wiring for power and Local BUS communications. These wires protrude through the centre of the mounting bracket which is fixed in place with screws. The main DVC enclosure is pushed onto the plastic plate and rotated until it is locked in place. This simultaneously connects the external power and communications terminals through to the electronics within the enclosure by using studs that wipe over a flexible portion of the terminals. A spring-like action ensures good contact between the terminals and the studs.

Throughout the DVC extensive use is made of two discrete logic families. The DVC uses two different supply rails for digital electronics, namely, 3.3V and 5V. Because of this mixed supply arrangement one of the logic families must be able to interface between the different operating levels effectively.

One of the logic families used is a standard advance CMOS logic series, 74ACxxx. Such devices can be powered by either the 3.3V power rail or the 5V power rail but devices powered by the 3.3V rail cannot be driven by devices powered by the 5V rail. This family is selected as high speed logic is necessary in the DVC and only small propagation delays can be tolerated in many areas. The 74ACxxx series is effective and is also readily available.

Another logic family used is the 74LCxxxx family which is specifically designed to be powered between 2.0V and 3.6V. These fast devices with small propagation delays have 5V tolerant inputs. Because the inputs can tolerate 5V signals these devices are particularly used where voltage level-conversion can be a problem.

To more fully describe the construction and operation of the DVC reference will now be made to the camera software operation and design.

As mentioned above, the DVC has three on-board microprocessors each of which performs a different function. While the same functionality could be achieved using one microprocessor this would complicate hardware and software design. As disclosed above, the three processors are a micro-controller in the image capture block 12, the dsp in the main processing block 11 and a

micro-controller in the communications block 10 for communications and controlling the down-loadable code.

The major components of the image capture micro-controller PIC code are a frame controller and an asynchronous serial communications to I²C communications converter.

Referring to Figure 2 the frame controller controls multiplexing of the image capture memory and banking of this memory.

Because the video memory can be accessed by the dsp 17 or the sensor 18 it is necessary to ensure that no accesses take place while the multiplexer 20 is switched to the wrong device. This is accomplished by using a token that is passed back and forth over the token interrupt lines 21 between the image capture controller 19 and the dsp 17. When the dsp 17 has the token the controller 19 switches the video memory to the dsp. When the dsp is finished with the imaging memory it returns the token to the controller 19. The controller 19 then waits until the next frame is due to be read out before switching the memory to the sensor 18. The sensor writes the frame and when it is finished the controller 19 sends the token back to the dsp.

The frame controller module on the controller 19 has several modes to put the image into the video memory in

different ways. The first way (default) is to simply put the image into bank B1 all the time. The second mode puts the image into bank B2 all the time. The third alternates the image between the two banks B1 and B2. The final mode captures two images in quick succession and puts them into bank B1 and B2.

The frame controller functionality could be implemented entirely in hardware.

To enable access to all the functionality and configuration control registers of the image sensor 18 a serial communications link is employed. In the preferred form the serial communication standard and protocol used (inter-integrated circuit - I^2C) is not available as a hardware port on the dsp 17. To communicate with this protocol it is necessary to implement the port in the software. This is achieved using the micro-controller of the image capture controller 19. The dsp communicates with the UART 22 of the controller 19 using standard asynchronous serial communications 23. The controller 19 takes the commands and data and passes them onto the image sensor 18. There are also commands from the dsp 17 to change settings and modes on the controller 19.

In an alternative arrangement the I^2C communications can be carried out directly from the dsp 17. This can be either a hardware I^2C port or a software implementation.

A sensor driver software module controls communications and interactions with the image capture controller 19 and the image sensor 18. The sensor driver software module has buffers for storing messages to go to the controller 19 and routines for processing messages returned. It also implements the sending and receiving of the video memory token.

The dsp code includes an image processing controller which acts as a wrapper for the image processing routines. Instead of calling the image processing routines directly when a new image is available (token received from controller 19) the image processing controller is called. The image processing controller is responsible for updating any information needed by the image processing routines, updating information that is needed for the image header, starting and stopping watchdog supervisory routines and sending the video memory token back to the image capture controller 19.

As previously identified, image compression is needed on the DVC because the communications bandwidth is severely limited for imaging. In the preferred form of the invention the compression technique used is a transform coder based on symmetric wavelet transform. Alternatives include fractal compression, JPEG (a standard with code readily available). The wavelet technique is preferred because of the availability of multi-resolution wavelet decomposed data for use in other image processor

algorithms. In addition fast execution and better compression than JPEG is achievable.

The image compression algorithm used is one of a class called lossy transform coders. The basic steps involved are as follows:-

- Transform the image.
- Quantise and truncate the coefficients of the transform (the lossy step).
- Entropy code the quantised coefficients.

The JPEG standard for image compression uses this sort of scheme. The transform used in JPEG is the Discrete Cosine Transform. This transform is applied to eight-pixel-square blocks of the image. The quantisation is done by dividing each output coefficient of the transform by a corresponding value from a quantisation table. The divided coefficients are then rounded. During this step, for normal photographic images, many of the coefficients are rounded to zero. The quantised coefficients are then zigzag scanned (to keep similar values together) and run length coded. The run length coding output is finally Huffman encoded to produce the output data stream. This scheme gives fairly good quality up to compression ratios of around 20:1 (0.4 bits per pixel).

In accordance with the present invention the compression algorithm follows the same basic scheme with different

components at some steps. The steps involved are as follows:-

Global (not blockwise) wavelet decomposition of the complete image to four levels.

Each level of the decomposition has its coefficients quantised and rounded.

The decomposition is scanned in a quad tree order.

Run length encoding is performed on the quantised, scanned decomposition.

Huffman encoding of the run length encoder output.

Scanning and transforming the whole image horizontally and then vertically produces an output with the complete image content represented by collections of sine and cosine waves (or other waveforms). An important fact for lossy image compression using transforms is that the human visual system is less sensitive to high spatial frequencies. The accuracy with which the high frequency components are produced (reducing the amount of information) does not adversely affect the appearance.

The wavelet transform separates out horizontal, vertical and diagonal detail into different images. The output of a transform level, or decomposition, is four blocks of

values. Each of approximately half the input dimensions and each representing one of horizontal detail, vertical detail, diagonal detail and a general approximation of the input image. This decomposition can be repeated, with the approximation image of the last level as input, to give successively smaller, more approximate images. The more approximate images have lower and lower spatial frequency components.

The quantisation is the "lossy" step to the compression. The quantiser maps the many possible outputs of the transform to a smaller number of possible "codes". The output of the transform step ideally results in many values that are close to zero. The quantiser will therefore map many of them to zero.

The scanning order seeks to maximise the effectiveness of the next step in the image compression. To do this a scanning order is needed that bunches up the similar values in the wavelet decomposed data together. This means that long runs of the same value are more likely to be obtained.

The run length coding replaces long "runs" of the same value with a single symbol and the number of times it occurs. Huffman coding simply maps the run length coding output codes to a set of variable bit length codes. In the variable bit length codes, the most common ones are the ones that have the fewest number of bits.

Motion detection in the DVC is based on inter-frame differences. The latest image is subtracted from the last. If there is a non-zero difference then this indicates that motion has occurred in the view between the two images. According to the present invention this differencing is carried out on a low resolution version of the viewed scene. This results in the differencing being more robust to small movements in the image (eg leaves fluttering in the wind).

Using a low resolution image means that less data has to be stored between frames. It also significantly reduces the computational load on the processor. This second advantage would not be possible if it was necessary to calculate the low resolution image before motion detection could be done. However, through the use of wavelet image compression as discussed above, it is possible to effectively get the low resolution image "for free". The wavelet transform produces a series of four successively lower resolution versions of the original image. Any one of these four levels could be selected to use for the motion detection.

Once it has been established that there is a difference between the two frames a determination can be made as to how significant is the difference. One method is to simply measure the absolute difference between the two frames. In the DVC the camera simply measures the mean

squared difference between the two frames. Once an estimate has been made it can be thresholded to make a motion/no motion decision.

The motion detection is applied to the image in ten or up to 24 different regions. The regions are user definable and each one generates a separate event if motion is detected.

The DVC according to the present invention can be readily developed to provide more sophisticated detection. For example, motion tracking could be employed and this may enable an alarm to be raised when a person goes in a door but not out.

The dsp code also provides image authentication watermarking. The purpose of this is referred to above.

According to a preferred arrangement the watermarking is carried out in two stages. First a digital signature of various data relating to the image is calculated. Secondly, part of this digital signature is used as a bit pattern that is embedded in the quantised wavelet coefficient output stream.

The digital signature is calculated as follows:-

The camera generates a private key/public key pair. The public key is sent to the head end PC. The

private key is stored in the flash memory of the camera and is never revealed.

The private key is used to calculate the digital signature of the following 264 bits of concatenated data:-

Checksum of the approximation image	128 bits
Image timestamp	32 bits
Image sequence number	16 bits
Camera serial number	48 bits
Site number	32 bits
Random Salt	8 bits

The standard ElGamal public key algorithm is used to generate the signature, but the DSA algorithm is used to calculate the public and private keys. A minimum key length of 528 bits for the public key and 132 bits for the private key is used.

This produces at least a 528 bit digital signature.

The first 96 bits of the digital signature is used as the watermark bit pattern. The remaining 432 bits of the digital signature are sent in the image header.

The 96 bits long bit pattern is repeated 6 times in the 512 x 384 sized image. By repeating the watermark it is possible to detect tampering in a region both by

detecting an invalid watermark and by detecting a region that disagrees with the majority.

To verify the watermark the 96 bits of the digital signature are first extracted and checked to see that the same 96 bits repeat across the image data. The digital signature is then reformed from the 432 bits in the image header and the 96 bits of watermark. The digital signature is then verified using the public key, which the front end PC has read from the camera.

The compressed image FIFO buffer is a part of the dsp code. The storing of compressed images is an important feature of the DVC.

In most imaging systems analogue video or digital, the camera simply sends out the image information as it is collected. This means that communications must be of sufficient bandwidth to cope with full rate real time images. It also implies that there is some means of buffering the imaging at the other end of the chain for review.

As mentioned previously, the DVC includes memory set aside to store several tens of images. This means that the DVC can be capturing images without sending any out and have past history available on demand should it be needed. Because there is not a constant stream of imaging information coming out, a low bandwidth digital

network can be used without fear of it being overloaded by imaging.

Unless adaptive, most lossy coders vary in performance with the complexity of the scene. All coder parameters being equal, a very complex scene (lots of detail) takes up more room than a less complex one. This means that the images stored in the DVC are assumed to be of variable length.

The requirements for image storage are:-

- Store as many images as is possible, delete them only when necessary.
- Be able to efficiently search the image buffer.
- Be able to send any of the stored images in any particular order.
- Images to be contiguous in memory to allow the use of DMA.
- Interface to put images in and read them out to be simple.

Figure 3 shows the organisation of the memory used for the image buffer 27. There is an array 28 of words which is statically declared. Space in array 28 is dynamically allocated to images as they are created. The images are indexed by an array of pointers 29 to image structures. The set of valid images are bounded by the oldest 0 to newest N images. These are indicated by pointers into

the image pointer array 29. Both the image data array 28 and the image pointer array 29 are circular in nature. The wrapping around is handled by image buffer functions. Images cannot be allocated over the end of the image data array 28. If they run off the end, they are moved to the front and the free space 30 at the end is left vacant.

The images in the image data array are stored compressed in a large, statically allocated array of integers (words). There is a header on the front of every image with the following format:-

```
|Image      Status|Size|Compression      settings      (3
words)|Request Number|Motion Status|Time Stamp|Data
```

Each of the header entries is one dsp word or 24-bits, except for the time stamp and compression settings. The image status is a state variable for a state machine that describes the behaviour of the image. Figure 4 shows the transitions of the image states and the functions that cause them. There are two sets of functions that modify the image status. the functions for building the image and the functions for sending the image.

The image size is the number of words in the data field. The time stamp is 32-bit UNIX time followed by a 16-bit sequence number.

The dsp code includes a host port communications controller. The host port of the dsp 17 is a 24-bit to 8-bit asynchronous parallel interface. It appears mapped in the I/O space of the processor in the communications block 10 as a collection of 8-bit registers. There are three transmit and receive registers. These represent the high, middle and low bytes of a single dsp word.

The processor in the communications block 10 sends "host commands" to the dsp 17 by writing 8-bit values to a register. These commands trigger a vectored interrupt in the dsp 17. There are four services the microprocessor in the communications block 10 can request from the dsp, these are:-

- Request a Local BUS packet from the dsp.
- Request an internal packet from the dsp.
- Request the dsp to get ready to receive a Local BUS packet.
- Request the dsp to get ready to receive an internal packet.

The host port communications controller software module on the dsp uses the dsp's Direct Memory Access controller to send the data to the host port. This saves on core processing and BUS usage.

The DVC maintains a real time clock which is used to add a time stamp to the images as they are stored. This clock can be updated from the Local BUS.

The dsp code includes message buffering. In addition to sending out images over the host port the DVC also sends and receives other Local BUS and internal packets. The message buffering simply queues up the messages to be sent or the messages received until they can be sent or processed respectively.

Finally, the dsp includes as a major component supervisory routines. The DVC runs several supervisory routines.

There is a check for the integrity of the code, a check for the correct functioning of memory and three watchdog routines. The watchdog routines check for hanging of the image processing, the sensor communications and the frame capture. If any problems are detected they are reported to the Local BUS either directly if the dsp 17 is capable, or indirectly by a termination mode via the controller of the communications block 10.

Major parts of the code of the controller of the communications block 10 are a dsp controller, Local BUS port controller and supervisory routines.

The dsp controller controls all of the interfacing to the dsp 17. It controls such things as the down-loading of code to the dsp, the hang supervision (watchdog) and the sending and receiving of packets from and to the dsp (including images).

The dsp code is down-loaded to the dsp whenever the dsp is reset. The bootstrap code on the dsp waits for the length and address of the code to be sent to the host port and then captures the code.

The Local BUS port controller is a serial communications controller peripheral (SCC) built into the same chip as the microprocessor of the communications block 10. The SCC handles much of the low level detail of the Local BUS communications. The Local BUS controller code handles the higher level part of the Local BUS protocol and processes the messages destined for the communication block 10 microprocessor.

The dsp has no hardware watchdog for hang prevention. Instead, the microprocessor of the communications block 10 functions as a watchdog. The dsp must send interrupts to the microprocessor periodically. If the microprocessor goes for a certain time without receiving an interrupt from the dsp it assumes the dsp has hung. It then checks the host port for a termination code from the dsp and this is used to help determine what has gone wrong with the dsp.

The communications block microprocessor supervisory routines also run some self-monitoring of the code integrity and memory functionality.

To more fully describe the invention reference will now be made to the circuitry of a form of the DVC. In doing so reference will be made to Figures 5-7 of the drawings.

Interfacing between the dsp and the APS is a microchip U25 (see Figure 6).

Referring to Figure 5, external connections for power and communications are made using a 6-way connector, P1, having the following pin designations:-

Pin 1:	COMA
Pin 2:	COMB
Pin 3:	OV (GND)
Pin 4:	OV (GND)
Pin 5:	POWER IN
Pin 6:	POWER IN

An external DC power source is in the range 10-30vdc and with a capacity of at least 3W. The positive potential connects to pins 5 and 6 and the return (OV) connects to pins 3 and 4.

External communications is a 2-wire, synchronous, BIT-orientated (FMO) protocol over a transformer-coupled RS-485 LAN referred to as the Local BUS. Connections are made to pins 1 and 2.

Transient voltage suppressors (Q1, Q4 and Q5) are placed between OV and the positive supply input, and between OV and each of the external communications lines.

The unregulated positive potential from the external power source is directed through a series, reverse-polarity, protection diode D1 and then an in-line fuse FS1. At this point the power is split in two directions: to a linear voltage regulator U38 that derives power for infra-red LEDs used for external illumination, and to the power supplies for the actual camera electronics. The physical split ensures that noise currents are not coupled from the infrared LEDs 31 into the camera electronics.

U38 is an adjustable linear voltage regulator with on/off control. The output voltage is set by the ratio of resistors R10 and R11. Capacitor C12 is provided on the output to ensure the stability of U38. Both the power to the LEDs 31 and its return line 32 are run separate from the main camera system. The return path 32 is through a series resistor R17 that sets the value of the current flow, and a FET (Q6) that provides the ability to turn the LEDs 31 on and off as required.

A simple Π -filter is used to clean the input power into the main camera electronics as well as to reduce conducted emissions from the camera back onto the external supply. The Π -filter consists of a small capacitor, C25, on the input side of an inductor L1, and slightly larger capacitors C10 and C11 on the other side of L1. This provides a protected, unregulated power rail referred to as V_{unreg} .

V_{unreg} is directed to the input (pin 2) of a Linear Technology buck-mode switching regulator U34 which is used to derive a 5V power rail for the digital components associated with the communications block 10 and image capture block 12. This 5V rail is referred to as V_d .

Regulator U34 has an on-chip bipolar NPN power switch. The emitter is available at pin 3 of the device. The collector is connected to input pin 2 via an on-chip series resistor that is monitored by a differential amplifier to provide a current-limiting function. Output pin 3 connects to the discrete regulating components: diode D2, inductor L2 and capacitor C13.

The switch operates at a constant rate of approximately 500kHz. During the switches on-time (about 1 μ s) pin is driven to the input voltage, V_{unreg} and the (linearly increasing) current flow through the inductor L2 results in energy being stored in its magnetic field. During the

switches off-time (also about 1 μ s) the magnetic field of L2 collapses. The resulting reverse polarity drives the potential at pin 3 negative and this is clamped by the catch diode, D2. A closed circuit is set up with L2, D2 and C13, and so the energy is transferred (as a linearly decreasing current) to C13 and the V_d .

Feedback is provided by connecting V_d to pin 7 of regulator U34 which is a fixed voltage part, meaning that the resistor divider normally associated with feedback senses is included on the chip. Two additional functions are performed by the feedback pin. When the pin voltage (ie the +5V rail) drops below 1.7V, the current limit of the switch is reduced. Below 1V the switching frequency is also reduced.

V_d also connects to pin 4 of U34. This provides a bias to improve efficiency when operating at higher input voltages and light load currents. Having this pin connected to the regulated output voltage forces most of the internal circuitry to draw its operating current from the output rather than the input supply.

Pin 8 of U34 is the output of the internal error amplifier and the input of the peak switch current comparator. It is used for frequency compensation, and is connected to capacitor C34 for this, however, it can double as a current clamp or control loop override. This pin sits at about 1V for light loads and 2V at maximum

load. It can be driven to ground to shut off the regulator.

A boost pin is provided on the pin 1 of U34. This is used to provide a drive voltage, higher than the input voltage (V_{unreg}), to the internal NPN power switch. Adding this additional voltage allows the switch to saturate and voltage loss approximates that of a 0.3Ω FET structure, improving switching efficiency. External boost components are capacitor C16 and diode D3. The anode of D3 is connected to V_d and this generates a voltage across C16 nearly identical to the regulated output. Efficiency is not affected by the capacitor value.

V_d is directed to a linear regulator (U35) which provides +3.3V for digital circuitry in the main processing block 11, video buffer block 13 and sensor/buffer control block 14. Capacitor C14 on the output of U35 ensures proper stable operation. U35 is a low-dropout regulator particularly designed for regulating from a 5V supply. Power dissipation is minimised by using a fairly low input voltage source.

V_{unreg} is also directed to a linear regulator (U42) which provides +5V used for analogue circuitry associated with the APS (U100) - see Figure 7. Both the supply and return lines to regulator U42 are filtered to reduce interference between the analogue and digital circuits.

Each line has a series inductor (L3 and L4) and there are capacitors (C36 and C26) across the lines either side of the inductors. The output of U42 is stabilised by capacitor C27. This +5V power rail is referred to as V_a .

Crystal oscillator X2 is a clock source for the microprocessor U31. This is divided down internally to provide a clock for U31. Capacitors C4 and C5 provide parallel loading for the crystal.

U31 interfaces directly to a 2M BIT FLASH E²PROM (an INTEL SmartVoltage Boot Block E28F200B), U32, and a standard 1M BIT SRAM, U33. U31 provides two software-programmable chip select lines labelled /ROMCS (used for the FLASH E²PROM) and /RAMCS (used for the SRAM). These signals are generated by decoding the address lines A19 to A12 of U31.

The lower address boundary for /ROMCS lower is fixed at 00000h, however, the upper boundary for /ROMCS, and both the lower and upper boundaries for /RAMCS, are software programmable. After a hardware reset the boundary addresses are set as follows:

/ROMCS lower address boundary (fixed):	00000h;
/ROMCS upper address boundary:	FFFFFh;
/RAMCS lower address boundary:	FFFFFh;
/RAMCS upper address boundary:	FFFFFh.

Priority is always given to /ROMCS over /RAMCS, and so with the above settings the /RAMCS will never be asserted until the control registers for these lines are re-initialised with lower values.

The chip select lines /ROMCS and /RAMCS from the U31 are logically OR'ed with the memory request, /MREQ, output signal also available from the U31 to ensure that the memory devices are enabled only during actual memory access cycles. Gates from a 74AC32 (U30) are used to achieve these.

The FLASH E²PROM has a 16kB, write-protected, boot-block sector, physically located at the bottom address of the device. When a DVC is manufactured, this sector may initially be the only non-volatile memory containing programme code which, functionally, must at least provide the mechanism to down-load new programme code from over the Local BUS. This is, it must contain sufficient code to communicate over the Local BUS, to provide the ability to copy code to the SRAM, to run code from the SRAM, and be able to programme the FLASH E²PROM while running from SRAM.

The boot-block write-protect input of U32 connects to a general purpose port pin of the microprocessor U31. When the write-protect input is at Logic low the boot sector is locked, preventing programme and erase operations. Because U31 can control this input, and because there is

a second device (the SRAM) to run code from, it is possible to rewrite the boot-block code. Such a task must obviously be performed in a well controlled manner. Particularly because the SRAM is not battery backed, a task such as this could easily be interrupted and leave the DVC in an unusable state.

One feature of the FLASH E²PROM U32 is a flexible BUS interface which can connect to either an 8-BIT (BYTE-wide) or 16-BIT (WORD-wide) data BUS. The mode-select input pin on U32 is tied low to select the 8-BIT option to suit the U31. Because of this flexible BUS feature, the mode selected, and the pin configuration of the U32 to cater for this, the lowest-order address line from U31 (A0) connects to an input on U32 referred to as 'A₁/D15' - a multi-functional pin which is D15 when the WORD-wide mode is selected. This means each of the remaining address lines, although contiguous, are staggered, with, for example, A1 from U31 driving the A0 input and A2 driving the A1 input.

Another feature of U32 is a reset control input pin. The system reset is connected to this pin on U32. When a *logic low* is applied, the device is in a rest/deep powerdown mode, which puts the outputs at high impedance, resets the internal write state machine and draws minimum current. When a *logic low-to-high transition* occurs on this input, the devices default to a read array mode - this is a necessary condition for the CPU to be able to

read from the device after a system reset has occurred. A logic high signal keeps the devices in standard operation.

Programming and erasing U32 can be achieved using either 5V or 12V. Reading from the U32 requires either 3.3V or 5V. This assembly uses 5V for all three operations. A general purpose port pin of U31 controls the output of an adjustable output linear voltage regulator U37 which provides the 5V programming voltage serviced from V_{unreg} . The ratio set by resistors R9, R13 and R14 in the adjustable feed-back loop set the output at 5V. Capacitor C15 stabilises this power rail. When U37 is switched off, the output drops to 0V which ensures absolute protection of the stored code-data. A pull-up resistor R15 on the on/off control input of the U37 ensures that the device is off whenever the microprocessor U31 port pin is set as an input, such as after a system reset.

Referring to Figure 6, the DSP(U1) core incorporates a versatile memory interface that provides a "glueless" interface to a variety of memory types such as DRAM, SRAM and synchronous SRAM. The DVC provides the dsp U1 with up to twelve fast SRAM (15 OR 20ns access time). The SRAM chips are all 128x8 Kbyte (1Mbit) devices which are standard parts available from a variety of manufacturers. Therefore, up to 12Mbits can be fitted to this assembly for general use by the dsp.

Because dsp U1 uses a 24-BIT data BUS and the SRAM are only 8-BIT devices, the SRAM are arranged in banks four of three U2-U13. Each bank being 128Kbyte x 24 bit = 128k WORDS.

U1 provides four software programmable address attribute pins, AA0-3. These are tri-statable outputs with a programmable polarity that can be used as chip select signals or additional address lines. Address lines A0 to A16 connect directly to all the general purpose SRAM devices U2-U13. AA2 and AA3 are logically combined with the A17 address line to provide four active-low chip select signals to access the four banks of SRAM as follows:

```

/CS1 = AA2 + A17      :selects US, U3 and U4,
/CS2 = AA2 + /A17     :selects U5, U6 and U7,
/CS3 = AA3 + A17      :selects U8, U9 and U10,
/CS4 = AA3 + /A17     :selects U11, U12 and U13

```

As discussed above image data is transferred between the APS U100 (an 8-BIT device) and the dsp U1 using a dual-port video buffer. Only one of the devices has access to the buffer at a time. The APS can only write data into the buffer. The dsp can both read from and write too the buffer. Access to the buffer is physically controlled by U28, however the controlling instructions are given to U28 by the dsp U1.

The buffer consists of three 128Kbyte x 8 bit fast SRAM IC's (U14, U15 and U16) arranged as though they were a single 128Kbyte x 24 bit device. They share a common 17 bit address BUS (VA0-16) and common control lines (/VRD, /VOE and /VWE), and each device connects to a different byte of the 24-bit data BUS - U14 connects to the low-order byte, U15 to the middle-order byte, and U16 to the high-order byte.

The sources of the control lines /VRD, /VOE and /VWE are selected using a quad 2-input multiplexer U26 which selects the appropriate signals from either the dsp or APS interface depending on which has valid access selected at the time. The selection is controlled by a signal, /VB_ACCESS, from a general-purpose port of U28. This line has a pull-up resistor to ensure that after a system reset the dsp has, by default, valid access to the video buffer.

When /VB_ACCESS is low the dsp has valid access to the video buffer and the control signals selected through U26 are derived as follows:

/WE: dsp /WR

/OE: dsp /RD

/CS: dsp AA0

/G: dsp AA0

The signal /G is explained below.

When /VB_ACCESS is *high* the APS has valid access to the video buffer and the control signals selected through U17 are derived as follows:

/WE	:	/SENSOR_WR	:	derived in the pld
/OE	:	V _d	:	held inactive
/CS	:	0V	:	held active
/G	:	V _d	:	held inactive

Under control of signals (Q1, Q2 and Q3) from the pld (U27) image data from the APS is latched one byte at a time into three D-type latches (U17, U18 and U19). Once each latch contains a data byte the contents of all three latches are written simultaneously into the video buffer and the next three image data bytes begin to accumulate in the latches. The address to which the contents are written is also provided from U27. When the contents from the latches have written into the video buffer the address is incremented ready for the next 24 bit WORD of data.

The WORD written into the video buffer has the order of the bytes mixed up. The first byte of image data that is loaded into the latches is written to the high-order byte data BUS of the video buffer, and the third byte of image data is written to the low-order byte data BUS. This facilitates the manner in which the dsp U1 handles the

data processing otherwise the dsp would have to swap the low-order and high-order bytes as each WORD was read. This effectively saves some dsp processing time.

The dsp U1 interfaces to the video buffer data BUS via three bidirectional BUS transceivers U20, U21 and U22. The direction select pin of each transceiver is controlled directly by the /RD signal from the dsp. The pin of the transceivers is controlled by the /G signal which is derived from the AA0 signal from the dsp via the quad multiplexer U26 when access to the buffer is available to the dsp. The address that the dsp is accessing is available on the video buffer address BUS via two more transceivers (U23 and U24). The direction of these are fixed (because addressing is a one-way function) and the output enable pins of the transceivers are controlled by the same signal from the multiplexer U26 that controls the output enable pins of U20-22.

The maximum data space require to store a full image frame from the PB-159DM U100 (see Figure 7) is 64kx24 - exactly half the size of the video buffer. Therefore, the buffer is divided into two banks of 64kx24 for the sensor interface. A bank select control line, VB_BANK_SEL, from the U28 is directed to the pld U27 which generates the required signals to control this. The dsp U1 sees the video buffer as a contiguous space of 128kx24.

Micro-controller supervisory functions (primarily for microprocessor U31) are performed by U36, a MAX706. This controls the resetting of U31 and U32 when either the V_d power rail is considered insufficient to power the electronics or the internal (to U36) watchdog timer is not toggled adequately by the U31. A pull-down resistor R27 on the reset line output ensures that the reset function operates adequately when the V_d drops below about 1V.

The watchdog timer is reset by toggling a general purpose port pin of the U31 which connects to the input for the watch-dog timer. This line must be toggled at least once every second to prevent U36 applying a hardware reset pulse to the DVC system. A pull-down resistor R29 on the watchdog timer input ensures that this input can never float, which would disable the watchdog function (hence the watchdog function is permanently enabled).

The watchdog output connects to the master reset input pin of U36, so that if the output goes low, due to insufficient toggling of the watchdog input, a reset is applied to the electronics.

U36 also monitors the input dc power supply, V_{unreg} , via a voltage divider (R8 and R52), and generates a power fail signal when the potential of the external power source drops below approximately 10V. Resistors R8 and R52 connect to V_{unreg} via a series diode D4 to ensure proper

decoupling of the voltage divider by capacitor C17. Without D4 and C17, dips on the external power supply (for instance, due to other devices being plugged onto the external power source) could prematurely trigger a power fail indication. The power fail signal drives the /INT2 input of U31, is directed to the dsp /NMI and a dsp general purpose port, and connects to a general purpose input of U31 so that brown-out conditions can be recognised properly.

The DSP U1 core features a PLL clock oscillator which allows the dsp to operate at a high internal clock frequency using a low frequency clock input. The clock generation is composed of two main blocks:

- PLL that performs:
 - clock input division;
 - frequency multiplication;
 - skew elimination.
- Clock generator that performs:
 - low power division;
 - internal and external clock generation.

A 20MHz clock oscillator X1 is used as the external source for the PLL and clock generator. X1 connects to the dsp via a simple low pass filter (R1 and C2). This filter is included to reduce emissions from this line.

The (filtered) oscillator output is also directed to an inverter/buffer U44 which feeds it to the U28.

PLL characteristics are determined by a capacitor C1 connected to the 3.3V supply. The PLL can multiply the input frequency by any integer number between 1 and 4096. This "multiplication factor", MF, affects the value of the PLL capacitor C1.

The PLL is capable of eliminating skew between the external clock source and the internal clock phases and output clock pin, providing the multiplication factor is less than or equal to 4.

An internal register BIT (PEN) enables the PLL operation. When PEN is set, the PLL is enabled and the internal clocks will be derived from the PLL VCO output. PEN may be set or cleared by software control any time during chip operation.

During hardware reset PEN receives the value of an external pin, PINIT (shared with /NMI). In the DVC, the PLL is enabled during a hardware reset as described in the following section.

Signals from U31 and U36 to the dsp U1 are buffered by a bi-directional transceiver U25 with 5V tolerant inputs. This is because U31 and U36 are part of a 5V system, whereas the dsp U1 is a 3.3V device. U25 is set to work

in one direction only, as an input buffer to the dsp. Outputs connect to a pld, the dsp reset input, dsp interrupt inputs, and the dsp SCI input.

A reset signal for the dsp is generated from a general-purpose port of U31. This can be activated either under software control or as a result of a hardware reset on U31. A pull-down resistor R39 ensures this line is in the active lowstate following a reset of U31, when the general-purpose port is set by default as an input. This "system reset" signal is directed through an inverter U45 to the output enable pin of U25 so that, when activated, it is actually disabling the UF tri-stating all its outputs. The output from U25 that drives the dsp reset input has a pull-down resistor R36 on it so that the active reset is applied to the dsp. This line is also directed to an input of a pld.

When the outputs of U25 are enabled (reset condition inactive) the respective input for the output that provides the system reset is held high.

The power fail signal, /PFAIL, from U36 is directed to the dsp /NMI dual-function input through U25. This line also connects to a general purpose input of the dsp. During a reset condition, while the outputs of U25 are at high impedance, a pull-up resistor R35 on the /NMI input ensures that a logic *high* is applied so as to enable the PLL (the dual function of this input).

The dsp U1 includes boot-strap ROM (192-WORD) which allows a variety of options for down-loading programme code to the dsp. Code can be loaded to any program RAM segment through the host interface, an external ROM, or the SCI port.

Four external dual-function pins (MODA, MODB, MODC, MODD) are used to indicate the method of boot-strap loading to be used after a hardware reset is applied to the dsp. These pins are read during a hardware reset. Each line is set by either a pull-up or pull-down resistor when U30 is tri-stated during the reset condition. The value associated with each pin is latched into an operating mode register when the /RESET pin is negated. Thereafter, these pins operate as general interrupt inputs.

The boot-strap mode selected for the DVC is for programme code to be down-loaded into programme RAM from the host interface programmed to operate in a PCI target (slave) mode. To select this mode of operation, the mode-select pins are initially set at:

MODD:MODC:MODB:MODA = 0100.

The programme code is down-loaded by U31. The host PCI bootstrap code expects to read a 24-bit WORD specifying the number of programme words, a 24-bit WORD specifying

the address (within the dsp program code RAM) to start loading the program words and then a 24-bit WORD for each program word to be loaded. These are stored in contiguous program RAM memory locations (within the dsp). After receiving the program words, program execution begins at the same location where loading started.

Each of the pins used to select the boot-strap mode during a hardware reset becomes an input for an interrupt after the reset condition is removed. The interrupt inputs are active low, Schmitt trigger inputs that are internally synchronised to the dsp clock output (CLKOUT) pin. They are level-sensitive or negative edge triggered and maskable. The following is a summary of the connections of these pins on this assembly:

a) MODA(/I RQA)

Only used for mode select function - interrupt capability not used. Connects to 0V through R32 as part of boot-strap mode selection.

b) MODB(/IRQB)

Interrupt input from a general purpose port of U31. This line is buffered through U25. During a hardware reset, when U25 is disabled, a pull-down resistor (R38) is used to set the state required for boot-strap mode selection.

c) MODC(/IRQC)

Interrupt input from a general purpose port U28. This line is buffered through U25. During a hardware reset, when U25 is disabled, a pull-up resistor R37 is used to set the state required for boot-strap mode selection.

d) MODD(/IRQD)

Only used for mode select function - interrupt capability not used. Connects to 0V through R31 as part of boot-strap mode selection.

The external Local BUS interface uses the internal 85C30 Serial Communications Controller (SCC) channel of U31. This SCC connects directly to a high speed, differential, tri-state BUS/line transceiver/repeater U39 designed to meet the requirements of the EIA standard RS-485. BUS faults that cause excessive power dissipation within the device trigger a thermal shutdown circuit, which forces the drivers into a high impedance state and indicates the occurrence of a fault on an output pin.

Four lines connect between the U31 SCC and U39: TxD, RxD, /RTS and /CTS. /CTS connects to the line-fault indicator output of U39 and /RTS connects to the data-enable input. The /RTS output additionally drives a circuit to pulse an

LED D8 that is used to indicate when data is being transferred. The /RTS signal is directed through two inverters of U45. When /RTS is active high, the inverter output charges a capacitor C28 via a series resistor R6 and a diode D5 (which is also in series because it is forward biased). As C28 charges rapidly, the input of an open-collector inverter U43 is driven high. This inverter is connected to LED D8 via a series current-limiting resistor R4. In this state it forward biases the LED and sinks current from it, thus the LED is turned on. When /RTS is low the output of inverter U45 goes low and sinks current from C28 through series resistors R6 and R12 (D5, which is across R12, is reverse biased and so R12 now plays a part). The extra resistance of R12 (previously bypassed with D5) results in a slower discharge time compared with the charge time. As C28 discharges it will slowly reach a level where inverter U43 changes state and the LED is turned off.

U39 drives into the primary winding of a LAN isolating transformer, TF1, and the secondary winding connects through to connector P1. Series capacitors C8 and C9 on either side of TF1 prevent the core of the transformer from being saturated by dc-components of signals. Resistor R7 across the Local BUS lines is required to ensure the correct line impedance matching. R7 is not always necessary and so a mini-jumper link (P3) is provided in series with it to connect or disconnect it as required.

The SCC derives its clock from a crystal X3 connected across two of its port pins. There is a capacitor to 0V on each pin (C6 and C7) to provide the necessary parallel circuit for the crystal.

U31 provides two UARTs. One of these connects through an RS-232/TTL converter U41 to provide an external, general-purpose, RS-232 serial port at connector P2. For protection there are series current-limiting resistors (R50 and R51) in the external receive and transmit data lines and transient voltage suppressors (Q2 and Q3) between each line and 0V. This port is provided for development and/or additional future functionality - possibly for connecting a monitor to a DVC when it is installed to, for example, check lens focus and scene coverage.

The dsp and U31 are connected with a high-speed parallel BUS interface, as though the dsp is a device in the U31 I/O space. The dsp 'host' port connects directly to the U31 data BUS lines ZD0-ZD7, address BUS lines ZA0-ZA2, read and write control lines, and a chip select signal, ZSEL, derived using two 74AC32 OR-gates (U30) and an inverter (U45).

ZSEL is an active low signal generated by the following logical combination:

ZSEL=/IORQ+ZA6+/ZA7.

This puts the dsp into the U31 I/O space at locations xx40 through xx4F.

Although the U100 is a single-supply device, requiring only 5V to operate, two 5V supplies, V_a and V_d , are connected to various pins associated with either analogue or digital internal circuitry respectively, both via simple filters.

The digital 5V power rail, V_d , supplies the digital sections of the APS via a simple n-filter consisting of capacitors C112 and C100 and inductor L10.

The analogue 5V power rail V_a , is directed straight to most of the analogue sections of the sensor and various bias circuitry (a set of voltage dividers). It is decoupled with a large capacitor, C101, and two small capacitors, C102 and C103. These capacitors also form part of a n-filter, along with C113, C104 and inductor L11 which directs power to a pin on the APS which powers the actual pixel sites

The dsp U1 core provides a dedicated user-accessible test access port (TAP) that is fully compatible with the IEEE 1149.1 Standard Test Access Port and Boundary Scan Architecture. This proposed standard was developed under the sponsorship of the Joint Test Action Group (JTAG) and

the Test Technology Committee of IEEE because of problems associated with testing high density circuit boards. Implementation on U1 core supports circuit-board test strategies based on this standard.

Five dedicated pins (4 signal pins and 1 reset pin) are provided for the TAP. These interface on-chip test logic consisting of a 16-state controller and 3 test data registers. A boundary scan register links all device signal pins into a single shift register. The test logic is independent of the device system logic. This provides the following capabilities:

1. perform boundary scan operations to test circuit-board electrical continuity;
2. bypass the U1 core for a given circuit-board test by effectively reducing the boundary scan register to a single cell;
3. sample the U1 core based device system pins during operation and transparently shift out the result in the boundary scan register. Preload values to output pins prior to invoking the boundary scan;
4. disable the output drive to pins during circuit-board testing;

5. provide a means of accessing the emulation circuits and emulation controller to control a target system;
6. provide a means of entering the "Debug Mode" of operation;
7. query identification information - manufacturer, part number and version; and,
8. force test data onto the outputs of U1 core based device while replacing its with a single BIT register.

Connector P5 provides access to the 5 dedicated JTAG pins, which are:

a) Test Clock (TCK)

The test clock input pin is used to synchronise the test logic.

b) Test Mode Select (TMS)

The test mode select input is used to sequence the test controller's state machine. The TMS is sampled on the rising edge of TCK and it has an internal pull-up resistor.

c) Test Data Input (TDI)

Serial test instructions and data are received through the test data input. TDI is sampled on the

rising edge of TCK and has an internal pull-up resistor.

d) Test Data Output (TDO)

The test data output is a serial output for test instructions and data. TDO is tri-stateable and is actively driven in some controller states. TDO changes state on the falling edge of TCK.

e) Test Reset (/TRST)

The test reset input is used to a synchronously initialise the test controller. The /TRST input from the connector P5 has a pull-up resistor, R40, and is directed to the pld where it is combined with the system reset signal before being directed to the /TRST pin of the dsp. This is to ensure that if there is a system reset then /TRST is applied simultaneously to the dsp to ensure correct operation occurs. Otherwise, the external /TRST can be applied without affecting the system reset line.

U1 is also equipped with on-chip emulation circuitry, which provides a means of interacting with the U1 core and its peripherals non-intrusively so that a user may examine registers, memory or on-chip peripherals to facilitate hardware and software development. Circuits and dedicated pins are defined to avoid sacrificing any on-chip resource. The emulator circuitry is accessed by a JTAG instruction.

Because the emulation controller functionality is accessed through the JTAG port, there are no dedicated emulator pins for clock, data in and data out. However, in order to facilitate emulation specific functions, an additional pin, referred to as /DE (Debug Event) can be used. This pin is also available at connector P5. It is an active low, bidirectional, open-drain port and so has an external pull-up resistor R41 attached to it.

As an input, /DE provides a fast means of entering the debug mode of operation from an external command controller. As an output, /DE provides a fast means of acknowledging the entering debug mode of operation to an external command controller. Asserting this pin causes the DSP56300 core to finish the current instruction being executed, save the instruction pipeline information, enter the debug mode of operation, and wait for commands to be entered from the TDI line.

If /DE was used to enter the debug mode then /DE must be negated after the on-chip emulator responds with an acknowledge and before sending the first emulator command. The assertion of the /DE pin by the U1 core indicates that the dsp has entered the debug mode and is waiting for commands to be entered from the TDI line.

The dsp's serial communications interface (SCI) provides a full-duplex port for serial communications to other

dsp's, microprocessors, or peripherals such as modems. This interface uses 3 dedicated pins: transmit data (TXD=PE1: TXD is labelled PE1 on circuit diagrams because the default configuration is as a general purpose I/O pin on port 'E'), receive data (RXD=PE0), and SCI serial clock (SCLK=PE2). It supports industry-standard asynchronous BIT rates and protocols as well as high-speed synchronous data transmission greater than 8Mbps. The asynchronous protocols include a multi-drop mode for master/slave operation.

The U1 SCI is linked to a USART on the U28 microprocessor. Only the TxD and RxD lines are connected. The TxD line from the dsp is buffered by two inverters (U44) in series, combining to be a non-inverting buffer between the 3.3V dsp and the 5V U28. The RxD to the dsp is buffered through bidirectional transceiver U25 with 5V tolerant inputs, U25.

A binary switch SW1 sets the unit address for the DVC in the Local BUS address architecture. The switch has 16 positions to select from. These are decoded onto four lines which are directed to general purpose ports of the U1 where the states of the lines can be monitored. Each line has a pull-up resistor (R46-R49) and the common connection of the switch goes to 0V. The equivalent binary number of the switch position is reflected onto the four lines.

Each DVC is identified by a unique 48-BIT serial number provided by a DS2401 (U40). This device has only data and ground pins and derives its power directly from the data line. Two general purpose port pins of the U1 are used to interface to the U40. One pin acts as a transmitter and connects to the U40 data line through an open-collector buffer U43 with a pull-up resistor R18. The other pin acts as a receiver and connects to the data line through a second open-collector buffer U43 with a pull-up resistor R16.

LED D7, which is a processing indicator, is driven by a general purpose port pin of U1 via open-collector buffer U43 and series current limiting resistor R5.

A general purpose port of U1 connects to the dsp circuitry so that U1 can reset the dsp if necessary. As the port is set as an input after a hardware reset there is a pull-down resistor R39 on this line to ensure the line is in its active state in that instance.

A general purpose port pin of U1 is connected though to the dsp to send an interrupt request to the dsp if required. This line is active-low and has a pull-up resistor R26 to disable this interrupt after a hardware reset.

There is also an interrupt request line from the dsp. This active-low line goes directly to the /INT1 input of the U1 and also has a pull-up resistor R24.

A general-purpose port on U1 is connected to the 'common' terminal of a micro-switch that is provided for tamper detection. The port line has a pull-up resistor R19. If the micro-switch is closed, then the port line is pulled to 0V through the switch. If it is open then the port line is disconnected from 0V and R19 pulls it high.

The present invention thus provides a digital video camera with integrated intelligence. Image data can thus be at least partially processed on board the camera such that priority data (eg for an alarm or unusual situation) is determined at the camera and transmitted to the front end of a monitoring system.

There is also provided a monitoring system in which data can be processed at the surveillance camera which reduces "information overload" at the front end of the system. Such a system is hereinafter described.

The monitoring system makes use of high bandwidth communication networks and advanced microprocessors to integrate access control, security alarm monitoring, digital imaging (with the DVC) and a digital intercom into one system.

The system has a command centre 40 which uses client server architecture. The server processes can be distributed over several machines or over a geographical area to implement a large system.

There may be numerous workstations 41 connected to the command centre 40. Each workstation 41 may have full system functionality or a restricted subset of functionality defined by operator or workstation.

The command centre 40 server(s) also perform a server function for access controllers 42 such as our CARDAX (Reg. Trade Mark) controllers but the system is designed so that a network of access controllers can continue normal access control operation if there is a failure of the command centre server 40a.

The CARDAX access controller 42 is preferred in the system and is an intelligent, microprocessor-based controller capable of communicating with 16 field devices. These devices include magnetic stripe and proximity card readers 43/44, input and output panel 45, DVCs 46 according to the present invention and digital intercoms 47 (embedded in some of the card readers).

For access control systems the CARDAX controller can control up to eight doors. The basic unit can store access details for 2,500 cardholders, and can have additional flash memory added on an expansion board to

allow access details for up to 200,000 cardholders to be stored.

The CARDAX controller 42 performs alarm and event monitoring, and will also be capable of programmable logic control (PLC) type functions. It will buffer and route image and intercom information to the appropriate server or workstation PC. A CARDAX controller can be used as part of a distributed access control or alarm monitoring system.

The CARDAX controller has a number of networking options to communicate to other CARDAX controllers and to the command centre. The primary method of communication is through standard Ethernet via a twisted pair cable. In the event of failure of the Ethernet a secondary connection can be made using point to point protocol (PPP), this may be either a permanent or dial up connection. PPP may also be used as the primary communication medium between the CARDAX controller and the command centre. In this case a large distributed network may be built utilising several forms of public or private communications networks including PSTN (public switched telephone network), the Internet, and other public or private data networks.

Communication to the field devices is by a high speed (187Kb approx.) RS485 network. Event data, configuration, commands, digital images, digital audio

and operational software are transferred across this network.

The system includes magnetic strip and proximity card readers 44. These readers will read a card and pass the data back to the CARDAX controller 42. The controller commands the reader to turn two on board LED indicators on or off and to sound an audible tone from a built in speaker.

According to a preferred form the system includes "plus" card readers 44 which have similar card reading functionality to a basic card reader 43. The reader includes a keypad with numeric keys and function keys, a 128x32-pixel LCD display and digital intercom.

The plus reader 44 in association with the controller 42 can implement a range of functionality:-

- Advanced access control prompting a user to enter a PIN.
- Intercom call (via function key) and intercom operation.
- System changes, via menus and function keys such as arming or disarming alarms, locking or unlocking doors and temporary configuration changes.

- Display of written messages in a variety of character sets to identified individuals after identifying himself or herself at the reader.
- Playing of recorded messages to people after they have identified themselves at the reader (by presenting their card).
- Playing of a pre-recorded sound track triggered by some system event (eg evacuation instructions following a fire alarm).
- Display of date and time.
- Display of welcome messages or advertising.

The system can include a range of input and output devices 45. All inputs are four state inputs that have two main states open/closed and include two extra states to detect faults and or tampering with detectors. Outputs are implemented as relays with dry contacts.

There are a range of options in I/O boards; 15 inputs and 5 outputs, 20 output and a unit with 2 outputs and 6 inputs (strike controller) has the I/O for the control of 2 doors. It also includes a communications repeater/isolator 48 that will protect the system from attack (attempts to violate security by shorting or feeding signals into the system cables) when units are mounted in external or public areas.

The DVC 46 connects to the same communications wiring as the card readers 43/44 and other field devices. An operator can view images from any camera 46 in the system on the operator workstation (PC) and configure the system to store several images from a given camera when a certain alarm occurs. As disclosed earlier the DVC is able to detect motion and generate an alarm as well as indelibly stamp images with a digital "watermark" to prevent the images being tampered with.

The system allows an operator at a PC workstation 41 to watch the images from the DVC as they are acquired by the camera with a frame up date rate of around 2 frames per second (for a 512 by 384 image compressed by around 35:1). If the images are smaller or compressed more a higher frame rate will be achieved.

The operator may pause on a particular frame and have the system verify the watermark of that image to be assured that the image came from a given camera and has not been altered or tampered with.

The occurrence of certain pre-selected events can trigger the transmission of stored images to a PC server such that the PC server will have recorded a number of images from before the event and a number of images from after the event. This means that the system does not record images to a hard disk continuously, but only stores

sequences of images around the time of events that are of interest.

Any event or alarm (or logical combination of events or alarms) that occur on any one of the controllers in the system may trigger recording of images. Examples of typical types of event are:-

- a person being granted access at a door with a card reader;
- a person being denied access at a door with a card reader;
- the contact on a reed switch opening to indicate that a door has opened;
- the contact on a reed switch opening to indicate that a window has opened;
- motion detected by the camera;
- motion detected by a passive infrared (PIR) detector;
- a glass break detector detecting a glass break.

An alarm is a type of event that is deemed to be out of the normal. The controller stores configuration to allow it to determine whether an event is an alarm based on the following criteria:-

- the source of the event (eg the specific contact that was being monitored);
- the type of the event (eg contact opened or contact closed);

- the time and date (eg a certain event is only an alarm outside normal working hours);
- whether the source and type of event is in a group of alarm points that is currently armed or disarmed. The group of alarm points may correspond to a geographical grouping of alarm points (eg alarm points for a specific area of a building) or may correspond to a logical grouping of alarm points (eg all the reed switches monitoring windows around the outside of a building). The group of alarm points may be armed or disarmed manually (from a card reader or operator workstation PC), automatically (when the area that the alarm points monitor is occupied or unoccupied) or automatically by time and date. An alarm point in the group of alarm points only generates an alarm when the group of alarm points is armed.

Hence the recording images to hard disk can also be controlled based on the above criteria.

The details of which event triggered the image recording is included with the data for the sequence of images when a sequence of images is sent to the PC server 40a for recording to hard disk. When saving the sequence of images to the hard disk the PC server 40a uses this information to link the record of the event in the systems event log database to the sequence of images. This allows the operator, when viewing the system's event

log on screen to navigate directly to the sequence of images and play the images back. For example, when acknowledging an alarm in the system, an operator can request to view any images associated with that alarm and immediately play back several seconds worth of images that were recorded around the time of the alarm occurrence.

Compressed images are transmitted to the controller 42 via the polled RS485 communications link from the DVC. The controller 42 then forwards the images to a PC 40a via an Ethernet network.

The camera 46 only sends images to the controller 42 when the controller has requested either (i) continuous live images, or (ii) event triggered images.

In the case of continuous live images, these are transmitted by the controller using UDP/IP over the Ethernet network directly to the PC workstation 41 at which the DVC 46 is being viewed.

In the case of event triggered images the images are transmitted by the controller using TCP/P over the Ethernet network to the PC server 40a for recording to hard disk.

All communications links are encrypted to prevent monitoring of these links, or injection of messages into the system by unauthorised parties.

The controller 42 buffers several seconds of images when transmitting live images to a PC workstation 41 to ensure that these images do not have to be re-sent by the DVC if they are needed for storing to hard disk on the PC server 40a.

The system allows an operator at a PC workstation 41 to switch on the intercom in any card reader 47 in the system that includes an intercom. The operator may then either listen to what is happening at the card reader or may talk to someone standing by the card reader.

A person at the card reader may call the operator by pressing a designated function key on the card reader that includes the intercom. This will generate an event that will be reported at the operator workstation. The operator may then accept the call by switching on the intercom.

The system may also be configured to use a simple switch or push button as an intercom call button.

While the operator is talking to or listening to the intercom the conversation may be optionally recorded to the hard disk of the workstation PC 41.

The occurrence of certain events can trigger the transmission of compressed audio data to a PC server 40a such that the PC server will have recorded a number of seconds of sound after the event. This means that the system does not record sound from an intercom to the hard disk continuously, but only stores sound around the time of events that are of interest.

Any event or alarm that occurs on any one of the controllers in the system may trigger recording of sound from an intercom. Examples of typical types of event have been described above.

An alarm is a type of event that is deemed to be out of the normal. The controller 42 stores configuration to allow it to determine whether an event is an alarm based on the criteria outlined previously.

Hence the recording sound to hard disk can also be controlled based on the above criteria.

The indexing of recorded audio on hard disk and audio transmission and buffering apply as has been described above in connection with images.

The card reader will immediately play back any compressed audio data that it receives.

In the case of an operator listening or talking to an intercom, the audio data is transmitted between the controller and the PC workstation using UDP/IP over the Ethernet.

The system as described uses a personal computer as a central processing point for system set-up, monitoring and reporting (ie system host). This computer connects over a high-speed communications medium to many access controllers 42.

The controllers 42 monitor and control access to doors. Each controller can interface to a variety of devices via the RS485 link. Compressed digital data will, on request, be sent from the DVC 46 over the RS485 link, to the controller 42 and then over the high-speed communications medium, to the system host 40.

While most data processing is performed at the system host 40 the DVC will be able to analyse rapidly captured images and perform functions described above. This allows the DVC to be used both as a surveillance camera and as a replacement for a standard PIR motion detector.

When a DVC is performing motion detection, it will be able to generate an event to be logged by the system host in the same manner as a contact alarm from a traditional PIR detector.


The DVC is fully integrated into the system such that it can be software controlled from an operator workstation 41. An operator will be able to select any camera in the system and have images from that camera displayed on their screen. Software development will allow an operator to compare images of a person from the DVC with an image held in a database.

The system thus makes use of a high bandwidth communication network and advanced processors to integrate devices such as access controllers, alarm monitoring digital imaging (via the DVC) and digital intercom to provide a highly effective, versatile and efficient system which does not suffer from front end information overload typical of known systems.

PEC (NEW ZEALAND) LIMITED
By its Attorney
DON HOPKINS & ASSOCIATES

Per: 

334 139







FIG. 7.



F I G . 8 .




FIG. 1.




FIG. 2.




FIG. 3.

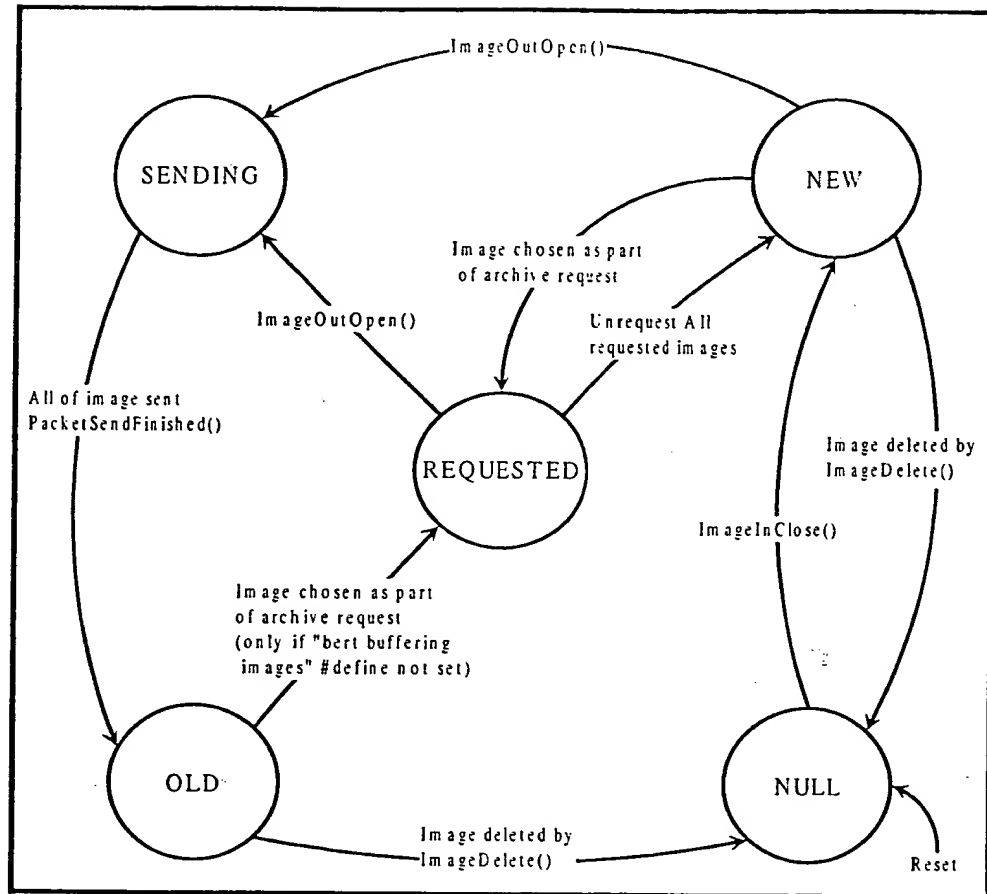


FIG. 4.

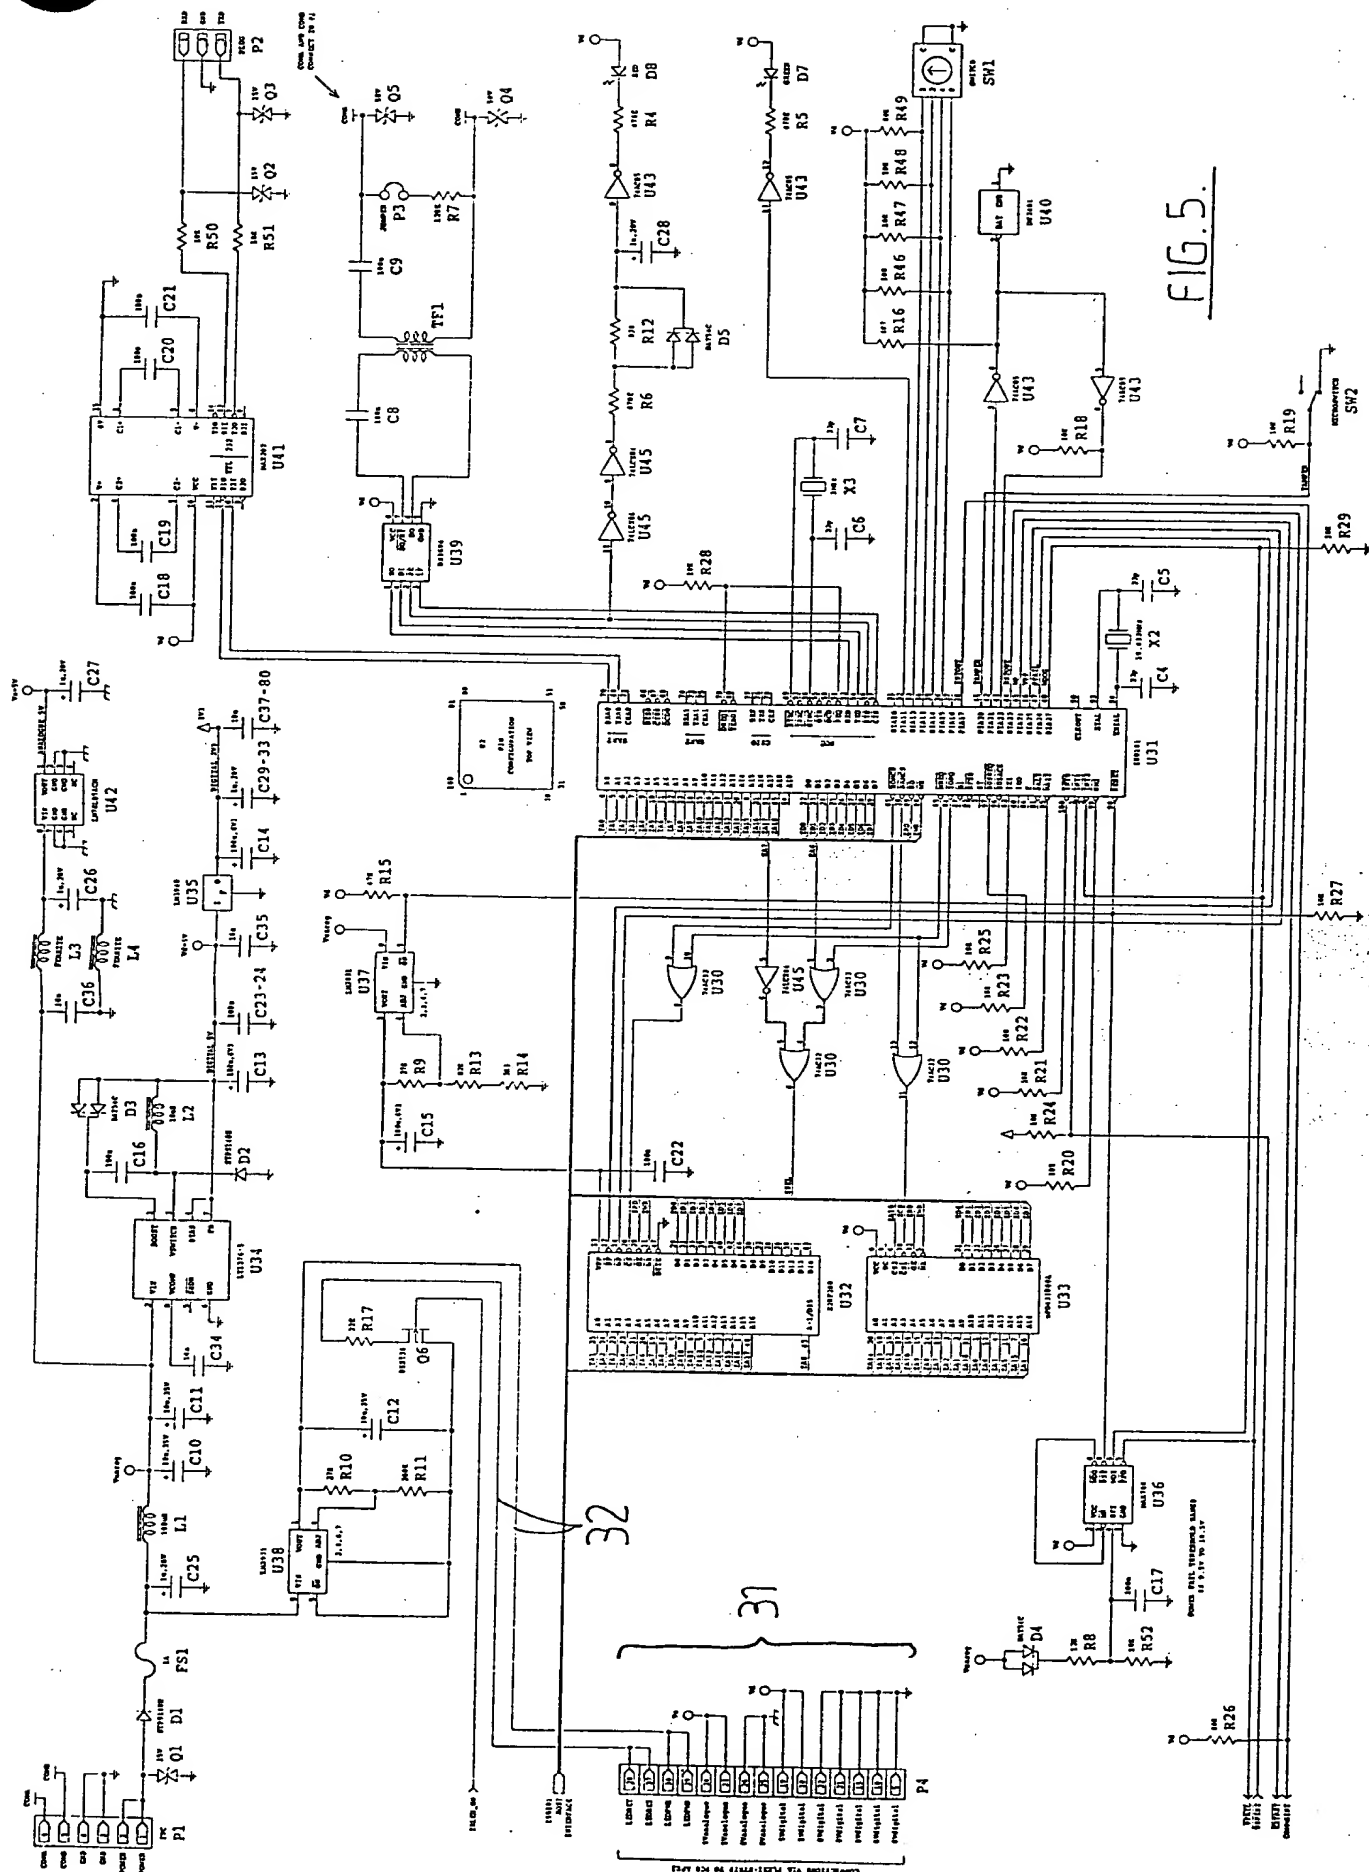


FIG. 5.

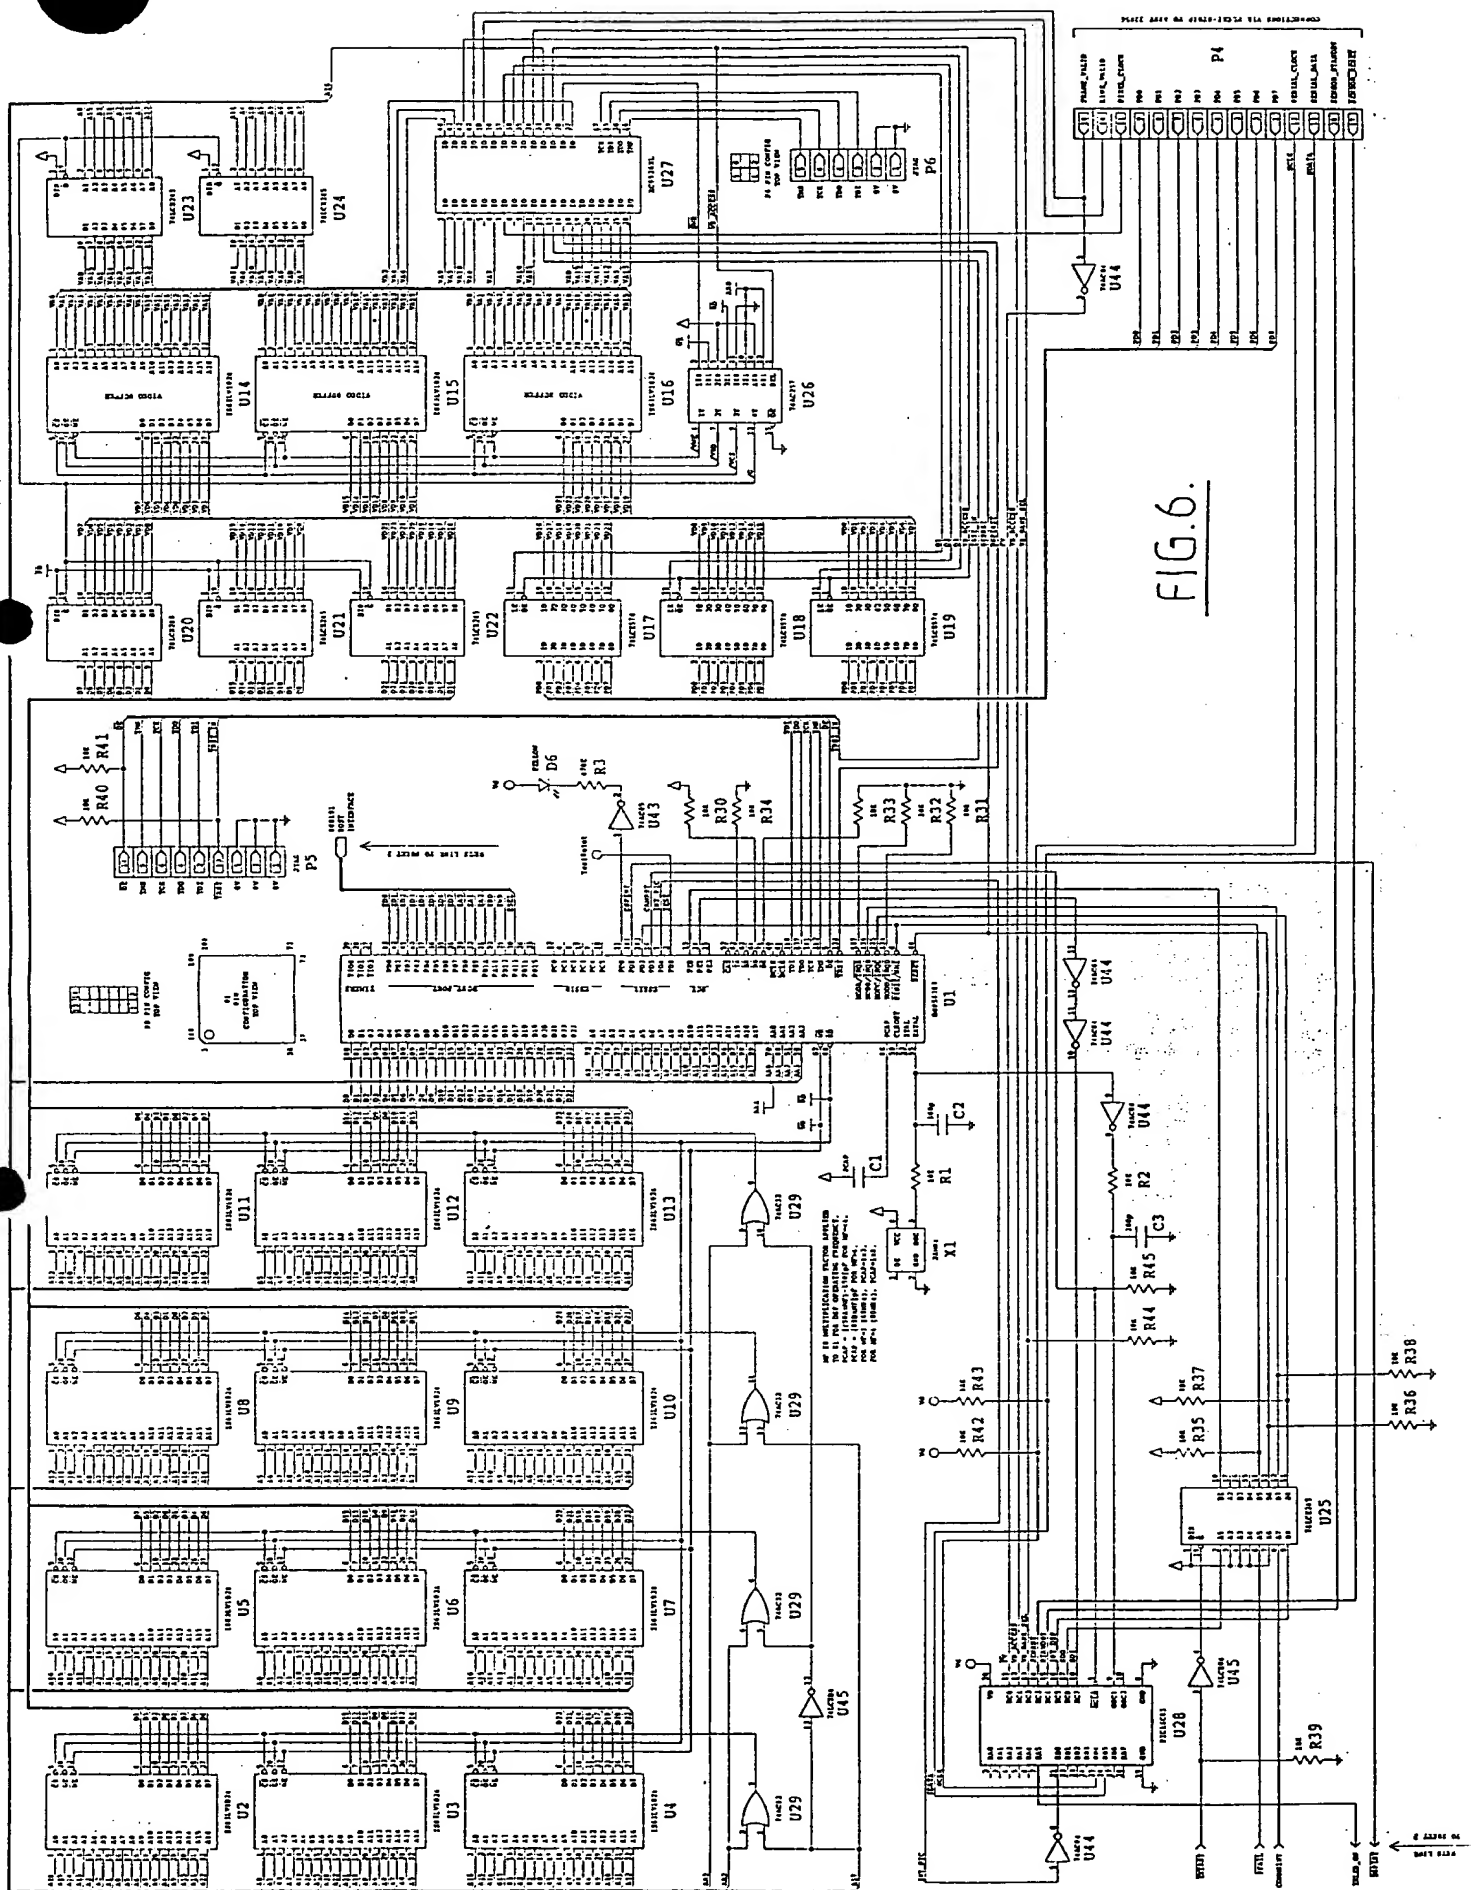


FIG. 6.




FIG 7.




FIG. 8.